

Permutation patterns: basic definitions and notation

David Bevan (The Open University)

Permutations, containment and avoidance

A permutation is considered to be simply an arrangement of the numbers $1, 2, \dots, n$ for some positive n . The length of permutation σ is denoted $|\sigma|$, and S_n or \mathfrak{S}_n is used for the set of all permutations of length n .

It is common to consider permutations graphically. Given a permutation $\sigma = \sigma(1) \dots \sigma(n)$, its *plot* consists of the points $(i, \sigma(i))$ in the Euclidean plane, for $i = 1, \dots, n$.




Figure 1: The plot of permutation 314592687 with a 1423 subpermutation marked

A permutation, or *pattern*, π is said to be *contained* in, or to be a *subpermutation* of, another permutation σ , written $\pi \leq \sigma$ or $\pi \preceq \sigma$, if σ has a (not necessarily contiguous) subsequence whose terms are order isomorphic to (i.e. have the same relative ordering as) π . From the graphical perspective, σ contains π if the plot of π results from erasing zero or more points from the plot of σ and then rescaling the axes appropriately. For example, 314592687 contains 1423 because the subsequence 4968 (among others) is ordered in the same way as 1423 (see Figure 1).

If σ does not contain π , we say that σ *avoids* π . For example, 314592687 avoids 3241 since it has no subsequence ordered in the same way as 3241.

If λ is a list of distinct integers, the *reduction* or *reduced form* of λ , denoted $\text{red}(\lambda)$, is the permutation obtained from λ by replacing its i -th smallest entry with i . For example, we have $\text{red}(4968) = 1423$. Thus, $\pi \leq \sigma$ if there is a subsequence λ of σ such that $\text{red}(\lambda) = \pi$.

Permutation structure

Given two permutations σ and τ with lengths k and ℓ respectively, their *direct sum* $\sigma \oplus \tau$ is the permutation of length $k + \ell$ consisting of σ followed by a shifted copy of τ :

$$(\sigma \oplus \tau)(i) = \begin{cases} \sigma(i) & \text{if } i \leq k, \\ k + \tau(i - k) & \text{if } k + 1 \leq i \leq k + \ell. \end{cases}$$

The *skew sum* $\sigma \ominus \tau$ is defined analogously. See Figure 2 for an illustration.

A permutation is called *sum indecomposable* if it cannot be expressed as the direct sum of two shorter permutations. A permutation is *skew indecomposable* if it cannot be expressed as the skew sum of two shorter permutations. Every permutation has a unique representation as the




Figure 2: The direct sum $2413 \oplus 4231$, the skew sum $2413 \ominus 4231$, and the layered permutation $21 \oplus 1 \oplus 321 \oplus 21$

direct sum of a sequence of sum indecomposable permutations, and also as the skew sum of a sequence of skew indecomposable permutations. If a permutation is the direct sum of a sequence of *decreasing* permutations, then we say that the permutation is *layered*. See Figure 2 for an example.

An *interval* of a permutation σ corresponds to a contiguous sequence of indices $a, a + 1, \dots, b$ such that the set of values $\{\sigma(i) : a \leq i \leq b\}$ is also contiguous. Graphically, an interval in a permutation is a square “box” that is not cut horizontally or vertically by any point not in it. Every permutation of length n has intervals of lengths 0, 1 and n . If a permutation σ has no other intervals, then σ is said to be *simple*.




Figure 3: The inflation $3142[123, 1, 21, 312] = 567198423$

Given a permutation $\sigma \in S_m$ and nonempty permutations τ_1, \dots, τ_m , the *inflation* of σ by τ_1, \dots, τ_m , denoted $\sigma[\tau_1, \dots, \tau_m]$, is the permutation obtained by replacing each entry $\sigma(i)$ of σ with an interval that is order isomorphic to τ_i . See Figure 3 for an illustration.

A simple permutation is thus a permutation that cannot be expressed as the inflation of a shorter permutation of length greater than 1. Conversely, every permutation except 1 is the inflation of a unique simple permutation of length at least 2.




Figure 4: The three left-to-right maxima and four right-to-left minima, and the two left-to-right minima and five right-to-left maxima, of a permutation

Sometimes we want to refer to the extremal points in a permutation. A value in a permutation is called a *left-to-right maximum* if it is larger than all the values to its left. *Left-to-right minima*, *right-to-left maxima* and *right-to-left minima* are defined analogously. See Figure 4 for an illustration.

Permutation statistics

An *ascent* in a permutation σ is a position i such that $\sigma(i) < \sigma(i+1)$. Similarly, a *descent* is a position i such that $\sigma(i) > \sigma(i+1)$. A pair of terms in a permutation σ such that $i < j$ and $\sigma(i) > \sigma(j)$ is called an *inversion*.

A permutation *statistic* is simply a map from the set of permutations to the non-negative integers. Classical statistics include the following:

- the number of descents
 $\text{des}(\sigma) = |\{i : \sigma(i) > \sigma(i+1)\}|$
- the number of inversions
 $\text{inv}(\sigma) = |\{(i, j) : i < j \text{ and } \sigma(i) > \sigma(j)\}|$
- the number of *excedances*
 $\text{exc}(\sigma) = |\{i : \sigma(i) > i\}|$
- the *major index*¹, the sum of the positions of the descents
 $\text{maj}(\sigma) = \sum_{\sigma(i) > \sigma(i+1)} i$

The statistics des and exc are equidistributed. That is, for all n and k , the number of permutations of length n with k descents is the same as the number of permutations of length n with k excedances. Furthermore, inv and maj also have the same distribution. Any permutation statistic that is distributed like des is said to be *Eulerian*, and a statistic that is distributed like inv is said to be *Mahonian*².

Classical permutation classes

The subpermutation relation is a partial order on the set of all permutations. A classical *permutation class*, sometimes called a *pattern class*, is a set of permutations closed downwards (a down-set) under this partial order. Thus, if σ is a member of a permutation class \mathcal{C} and τ is contained in σ , then it must be the case that τ is also a member of \mathcal{C} . From a graphical perspective, this means that erasing points from the plot of a permutation in \mathcal{C} always results in the plot of another permutation in \mathcal{C} when the axes are rescaled appropriately. It is common in the study of classical permutation classes to reserve the word “class” for sets of permutations closed under taking subpermutations.

It is natural to define a classical permutation class “negatively” by stating the minimal set of permutations that it avoids. This minimal forbidden set of patterns is known as the *basis* of the class. The class with basis B is denoted $\text{Av}(B)$, and $\text{Av}_n(B)$ or $S_n(B)$ is used for the set of permutations of length n in $\text{Av}(B)$. As a trivial example, $\text{Av}(\mathbf{21})$ is the class of increasing permutations (i.e. the identity permutation of each length). As another simple example, the class of **123**-avoiders, $\text{Av}(\mathbf{123})$, consists of those permutations that can be partitioned into two decreasing subsequences.

The basis of a permutation class is an antichain (a set of pairwise incomparable elements) under the containment order, and may be infinite. Classes for which the basis is finite are called *finitely based*, and those whose basis consists of a single permutation are called *principal* classes.

¹Named after Major Percy Alexander MacMahon.

²See footnote 1.

Non-classical patterns

Permutation patterns have been generalised in a variety of ways.

A *barred* pattern is specified by a permutation with some entries barred ($\bar{5}\bar{3}\bar{2}\bar{1}4$, for example). If $\hat{\pi}$ is a barred pattern, let π be the permutation obtained by removing all the bars in $\hat{\pi}$ (53214 in the example), and let π' be the permutation that is order isomorphic to the non-barred entries in $\hat{\pi}$ (312 in the example). An occurrence of barred pattern $\hat{\pi}$ in a permutation σ is then an occurrence of π' in σ that is not part of an occurrence of π in σ . Conversely, for σ to avoid $\hat{\pi}$, every occurrence in σ of π' must feature as part of an occurrence of π .

A *vincular* or *generalised* pattern specifies adjacency conditions. Two different notations are used. Traditionally, a vincular pattern is written as a permutation with dashes inserted between terms that need not be adjacent and no dashes between terms that must be adjacent. Alternatively, and perhaps preferably, terms that must be adjacent are underlined. For example, 314265 contains two occurrences of $\underline{231}4$ (or $2-31-4$) and a single occurrence of $\underline{231}4$ ($2-314$), but avoids $\underline{2314}$ ($23-14$).

A vincular pattern in which *all* the terms must occur contiguously is known as a *consecutive pattern*.

In a *bivincular* pattern, conditions are also placed on which terms must take adjacent *values*.

Classical, vincular and bivincular patterns are all example of the more general family of *mesh* patterns. Formally, a mesh pattern of length k is a pair (π, R) with $\pi \in S_k$ and $R \subseteq [0, k] \times [0, k]$ a set of pairs of integers. The elements of R identify the lower left corners of unit squares in the plot of π , which specify forbidden regions. Mesh pattern (π, R) is depicted by a figure consisting of the plot of π with the forbidden regions shaded. See Figure 5 for an example.




Figure 5: Mesh pattern $(3241, \{(0, 2), (1, 3), (1, 4), (4, 2), (4, 3)\})$

An occurrence of mesh pattern (π, R) in a permutation σ consists of an occurrence of the classical pattern π in σ such that no elements of σ occur in the shaded regions of the figure. A vincular pattern is thus a mesh pattern in which complete columns shaded.

Sets of permutations defined by avoiding barred, vincular, bivincular or mesh patterns that are not closed under taking subpermutations are known as *non-classical* permutation classes.

Growth rates

Given a permutation class \mathcal{C} , we use \mathcal{C}_n to denote the permutations of length n in \mathcal{C} . It is natural to ask how quickly the sequence $(|\mathcal{C}_n|)_{n=1}^{\infty}$ grows.

In proving the Stanley–Wilf Conjecture, Marcos and Tardos established that the growth of every classical permutation class except the class of all permutations is at most exponential.

Hence, the *upper growth rate* and *lower growth rate* of a class \mathcal{C} are defined to be

$$\overline{\text{gr}}(\mathcal{C}) = \limsup_{n \rightarrow \infty} |\mathcal{C}_n|^{1/n} \quad \text{and} \quad \underline{\text{gr}}(\mathcal{C}) = \liminf_{n \rightarrow \infty} |\mathcal{C}_n|^{1/n}.$$

The theorem of Marcos and Tardos states that $\overline{\text{gr}}(\mathcal{C})$ and $\underline{\text{gr}}(\mathcal{C})$ are both finite.

When $\overline{\text{gr}}(\mathcal{C}) = \underline{\text{gr}}(\mathcal{C})$, this quantity is called the *proper growth rate* (or just the *growth rate*) of \mathcal{C} and denoted $\text{gr}(\mathcal{C})$. Principal classes, those of the form $\text{Av}(\pi)$, are known to have proper growth rates. The growth rate of $\text{Av}(\pi)$ is sometimes known as the *Stanley–Wilf limit* of π and denoted $L(\pi)$. It is widely believed, though not yet proven, that every classical permutation class has a proper growth rate.

Wilf equivalence

Given two classes, \mathcal{C} and \mathcal{D} , one natural question is to determine whether they are equinumerous, i.e. $|\mathcal{C}_n| = |\mathcal{D}_n|$ for every n . Two permutation classes that are equinumerous are said to be *Wilf equivalent* and the equivalence classes are called *Wilf classes*. If principal classes $\text{Av}(\sigma)$ and $\text{Av}(\tau)$ are Wilf equivalent, we simply say that σ and τ are Wilf equivalent.

From the graphical perspective, it is clear that classes related by symmetries of the square are Wilf equivalent. Thus, for example, $\text{Av}(\mathbf{132})$, $\text{Av}(\mathbf{231})$, $\text{Av}(\mathbf{213})$ and $\text{Av}(\mathbf{312})$ are equinumerous. However, not all Wilf equivalences are a result of these symmetries. Indeed, as is well known, both $\text{Av}(\mathbf{123})$ and $\text{Av}(\mathbf{132})$ are counted by the Catalan numbers, so all permutations of length three are in the same Wilf class.

Generating functions

The ordinary *generating function* of a permutation class \mathcal{C} is defined to be the formal power series

$$C(z) = \sum_{n \geq 0} |\mathcal{C}_n| z^n = \sum_{\sigma \in \mathcal{C}} z^{|\sigma|}.$$

Thus, each permutation $\sigma \in \mathcal{C}$ makes a contribution of $z^{|\sigma|}$, the result being that, for each n , the coefficient of z^n is the number of permutations of length n . Clearly, two classes are Wilf-equivalent if their generating functions are identical.

A generating function is *rational* if it is the ratio of two polynomials. A generating function $F(z)$ is *algebraic* if it can be defined as the root of a polynomial equation. That is, there exists a bivariate polynomial $P(z, y)$ such that $P(z, F(z)) = 0$.

References

- [1] Miklós Bóna. *Combinatorics of Permutations*. Discrete Mathematics and its Applications. CRC Press, second edition, 2012.
- [2] Sergey Kitaev. *Patterns in Permutations and Words*. Springer, 2011.